



**Barcelona
Supercomputing
Center**

Centro Nacional de Supercomputación



EXCELENCIA
SEVERO
OCHOA

Increasing EC-Earth productivity & usability

Update on BSC activities

Miguel Castrillo



- New machine → Scalability results
- Take the most of the machine -> Performance?
 - ELPiN
 - IFS with XIOS I/O server
- Experiment productivity
 - Autosubmit: Visualization improvements and wrapper
 - Auto-EC-Earth: Branching strategy
 - ece2cmor3
 - Cylc proof of concept
- New developments
 - Stratospheric aerosols
 - T1279-ORCA12
 - Ensemble Kalman Filter
 - Nudging, atmospheric forcing set management
 - Initial conditions creation
- Evaluate changes and platforms → Reproducibility



EC-Earth model performance

- MareNostrum IV in operation since July 2017
- One of the first HPCs featuring new Intel Scalable Processors

	MareNostrum III	MareNostrum IV
Processor	Intel Xeon E5-2670 2.6 GHz	Intel Xeon Platinum 8160 2.1 GHz
#Cores per socket	8	24
#Sockets	2	2
Memory	32Gb DDR3-1600 2 GB/core	96Gb DDR4-2667 2 GB/core
Interconnection	Infiniband FDR10 10Gb	Intel Omni-Path 100Gb




MareNostrum III - 11,15 petaFLOPS




MareNostrum IV - 11,15 petaFLOPS


Scalability for EC-Earth trunk with default output configuration.



- ELPiN allows to find proper namelist parameters to **exclude land-only** processes in NEMO simulations
- Avoids NEMO to waste resources, **speeds up** simulations.
- Included in util/ folder in the project PRIMAVERA branch.



ORCA025 domain decomposed in 1287 sub-domains. 312 are land-only and therefore removed (24% of the total grid).
O. Tintó (BSC)

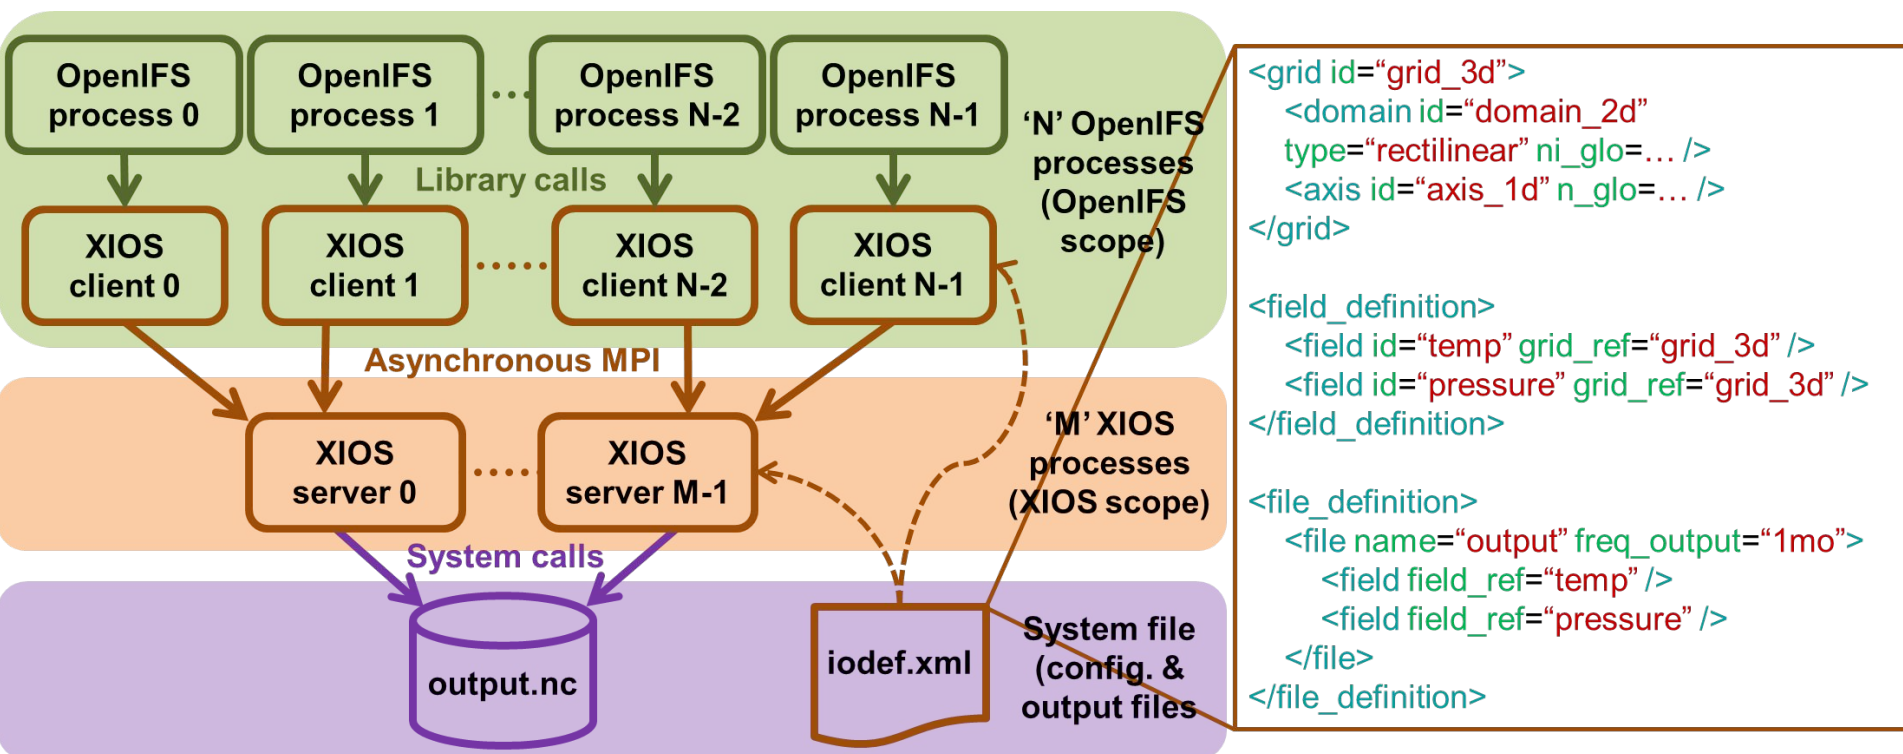


Impact of ELPiN on the NEMO model for an ORCA025-LIM3 simulation.
O. Tintó (BSC)

- Input/Output (I/O) management and post-processing increasing in **cost**, **complexity** and size
- IFS needs a **new approach** for EC-Earth scalability
- Goal: **XIOS** integration into the **OpenIFS** version that will be used in **EC-Earth**
 - Current development for IFS43r3 and future similar **OpenIFS** version
 - Three collaborating institutions:
 - BSC (Xavier Yepes and Mario Acosta)
 - ECMWF (Glenn Carver)
 - eScienceCenter (Gijs van den Oord)




- Furthermore, unlike the current solution (FullPos), XIOS writes in NetCDF format, which is a requirement of **CMIP6**





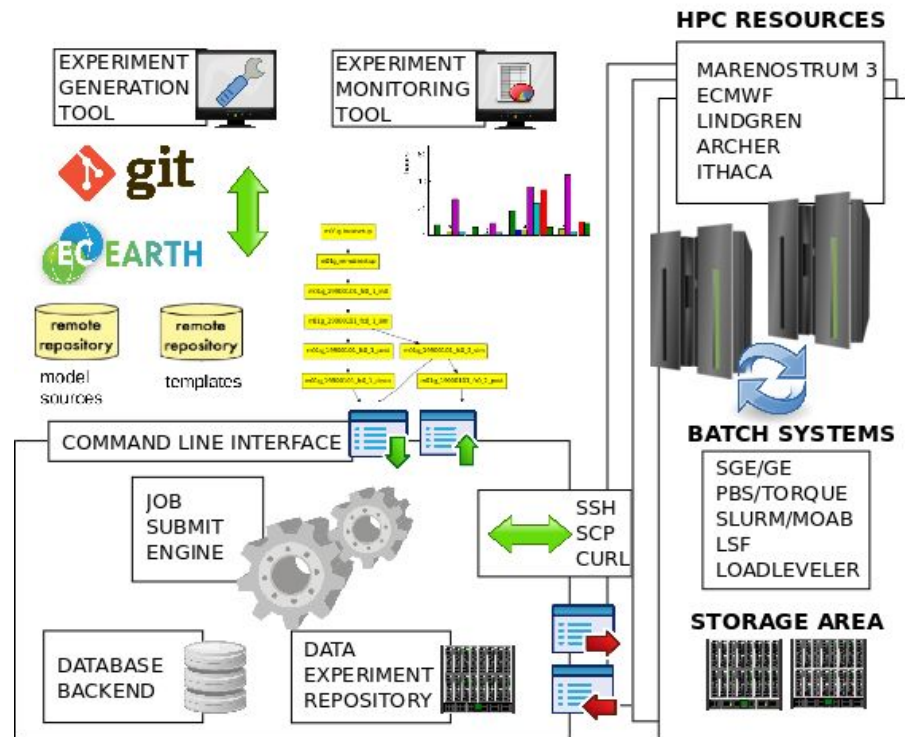
Experiment productivity (workflow managers)

- Workflow managers are **essential** to carry out production experiments in an **efficient** manner.
 - Deal with workflow **complexity**
 - Ensure **robustness & portability**
 - Usability** → Scientists more productive




- Autosubmit

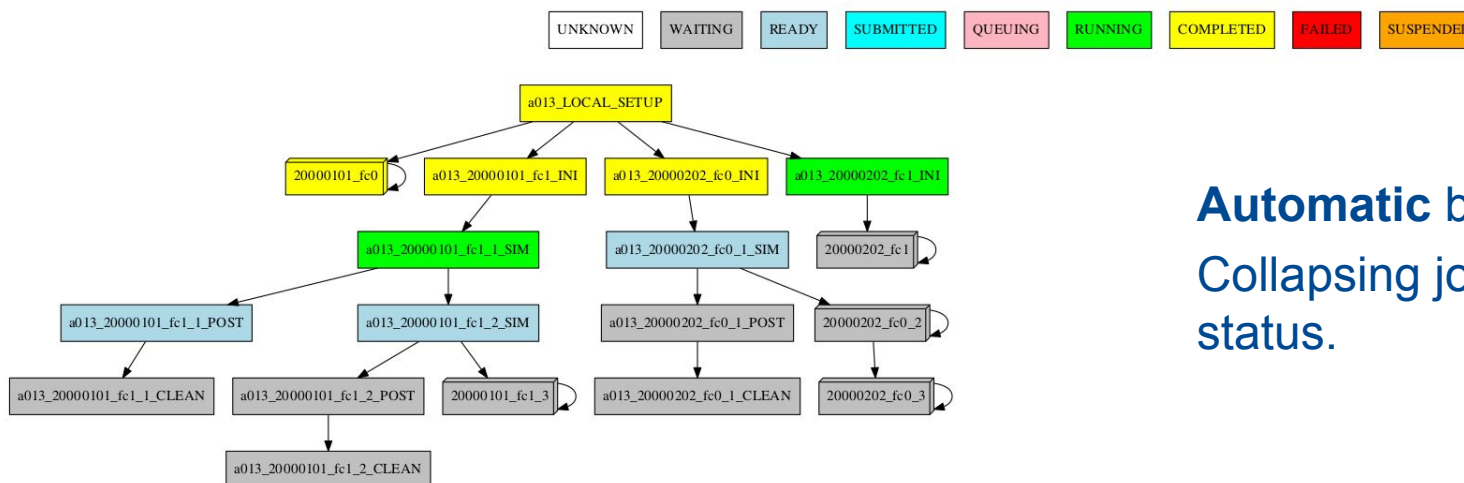
- A versatile tool to **manage** Weather and Climate **Experiments** in **diverse** Supercomputing Environments
- <https://pypi.python.org/pypi/autosubmit>



- Workflows are getting more and more **complex**
- Workflow managers are required to **improve** in order to **deal** with this **complexity**

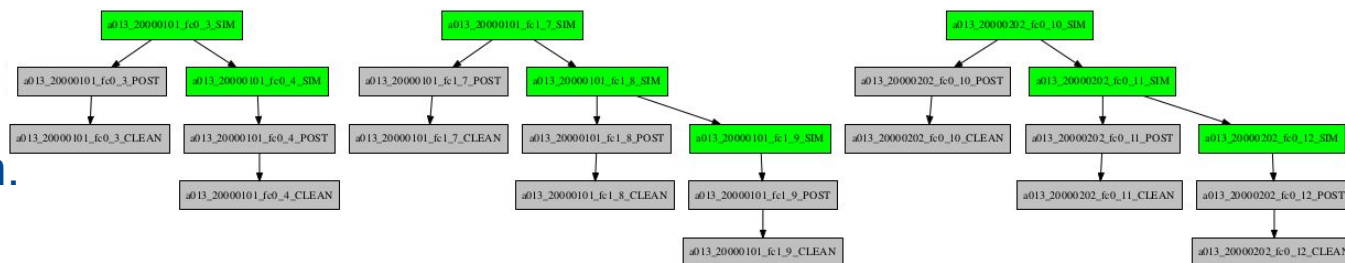


Improvements in the graph **visualization** of the workflow, by **grouping** jobs by date, member, chunk, split; or automatically.

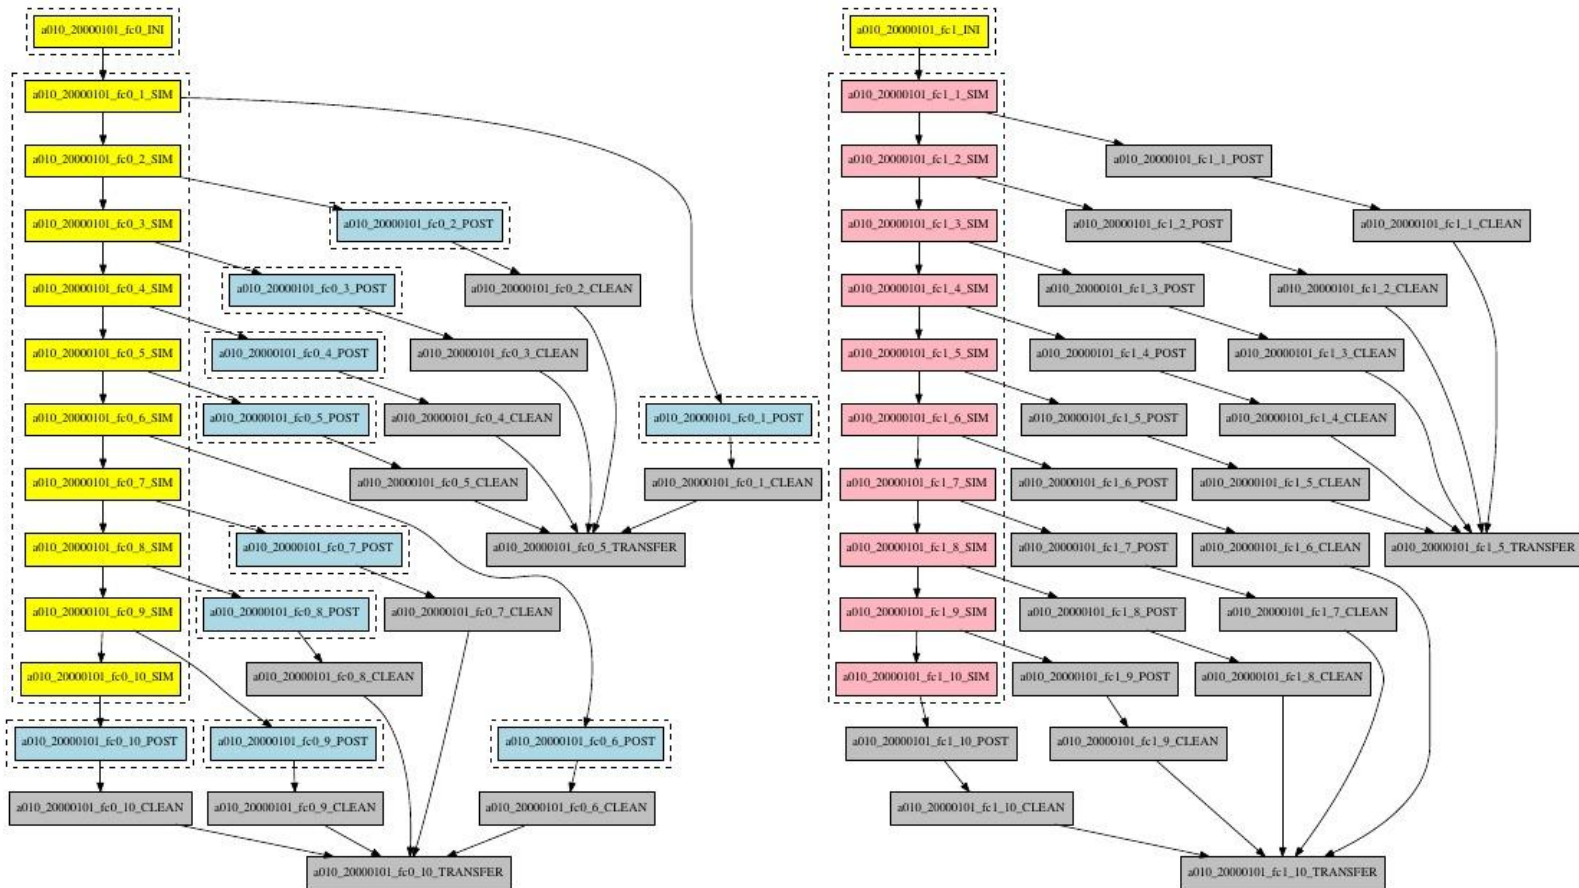


Automatic behavior:
Collapsing jobs sharing status.

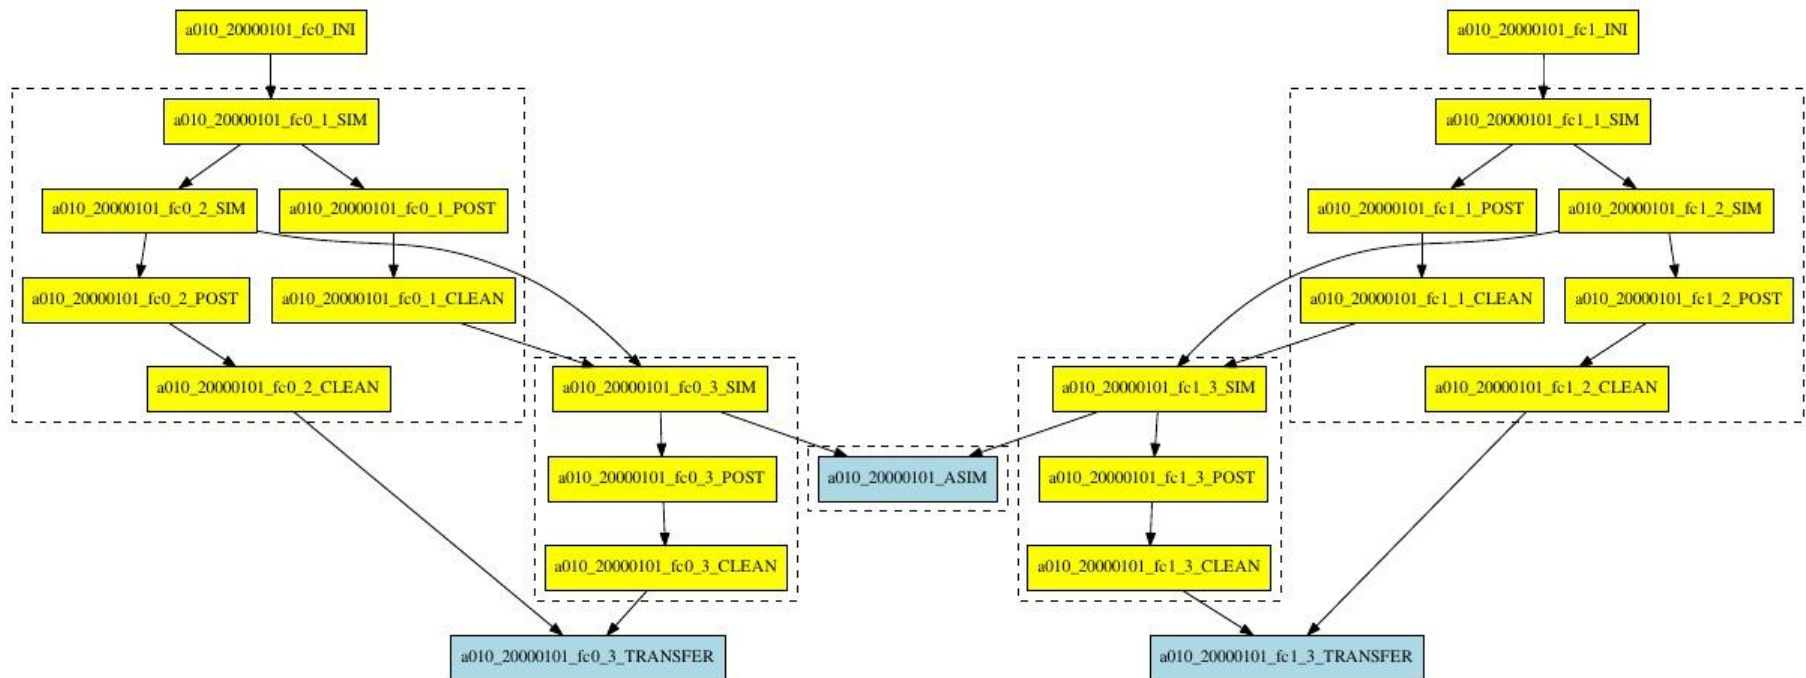
Hide groups:
Showing the more
relevant information.



Motivation: to **improve** the throughput by **reducing** queueing time by wrapping different jobs together.



Since last version it is also possible to wrap together jobs from different kind.



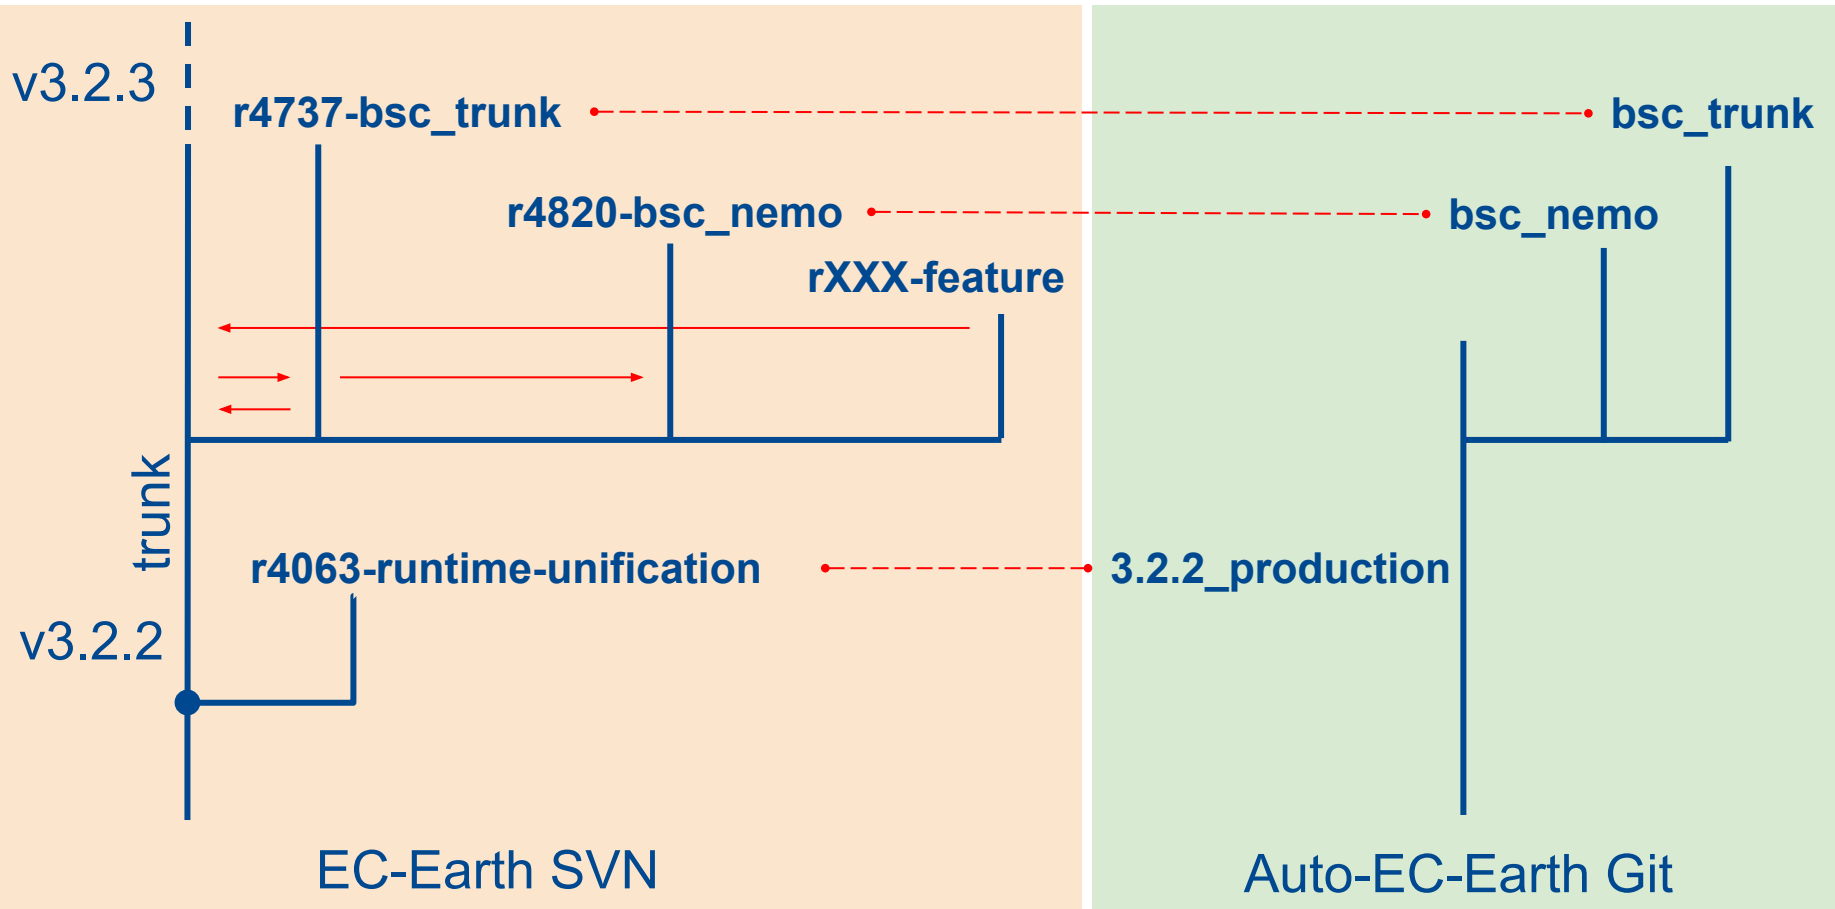
Auto-EC-Earth branching




Goal: To be able to **run** the last EC-Earth version, **use** new features, **merge** last developments smoothly

EC-Earth

Auto-EC-Earth





Every week: run a set of LR tests

bsc_trunk

auto-ecearth branch -> r4737-bsc_trunk

nemo only

ifs only

couple

cold start in CCA and in MN4

bsc_nemo

auto-ecearth branch -> r4820-bsc_nemo

nemo only

cold start in MN4

Every release: run a complete set of tests

nord3	CCA	MN4	resolution	type	details
t02c	t00u	t00q	T255L91-ORCA1L75-LIM3	coupled	start from restart
		t00v	T255L91-ORCA1L75-LIM3	coupled	ATM nudgin
		t011	T255L91-ORCA1L75-LIM3	coupled	sppt
	t00s	t00o	T255L91-ORCA1L75-LIM3	coupled	cold start
	t01d	t00z	ORCA1L75-LIM3	nemo	cold start
		t01j	ORCA1L75-LIM3	nemo	cold start ocean nudging
	t01e	t00r	T511L91-ORCA025L75-LIM3	coupled	start from restart
		t01o	ORCA025L75-LIM3	nemo	cold start
	t01b	t00y	T511L91	ifs	cold start
	t00t	t00p	T511L91-ORCA025L75-LIM3	coupled	cold start

- **Integration** of ece2cmor3 in Autosubmit/Auto-EC-Earth workflow
- Systematic **online** (=on HPC) **cmorization** of all experiments (PRIMAVERA/CMIP6 or “hybrid/local” configurations)
- Adapted “local” tables to handle specific configurations (model and pressure levels, frequencies, variables,...)

=> reduction of output size to transfer, standardization of ocean and atmospheric files, use of HPC capabilities for high memory operations, no need for extra post-processing, all tools have been adapted to work with CMORized files (facilitate the use of external data)

IFS grib (high freq)

Nemo raw NetCDF

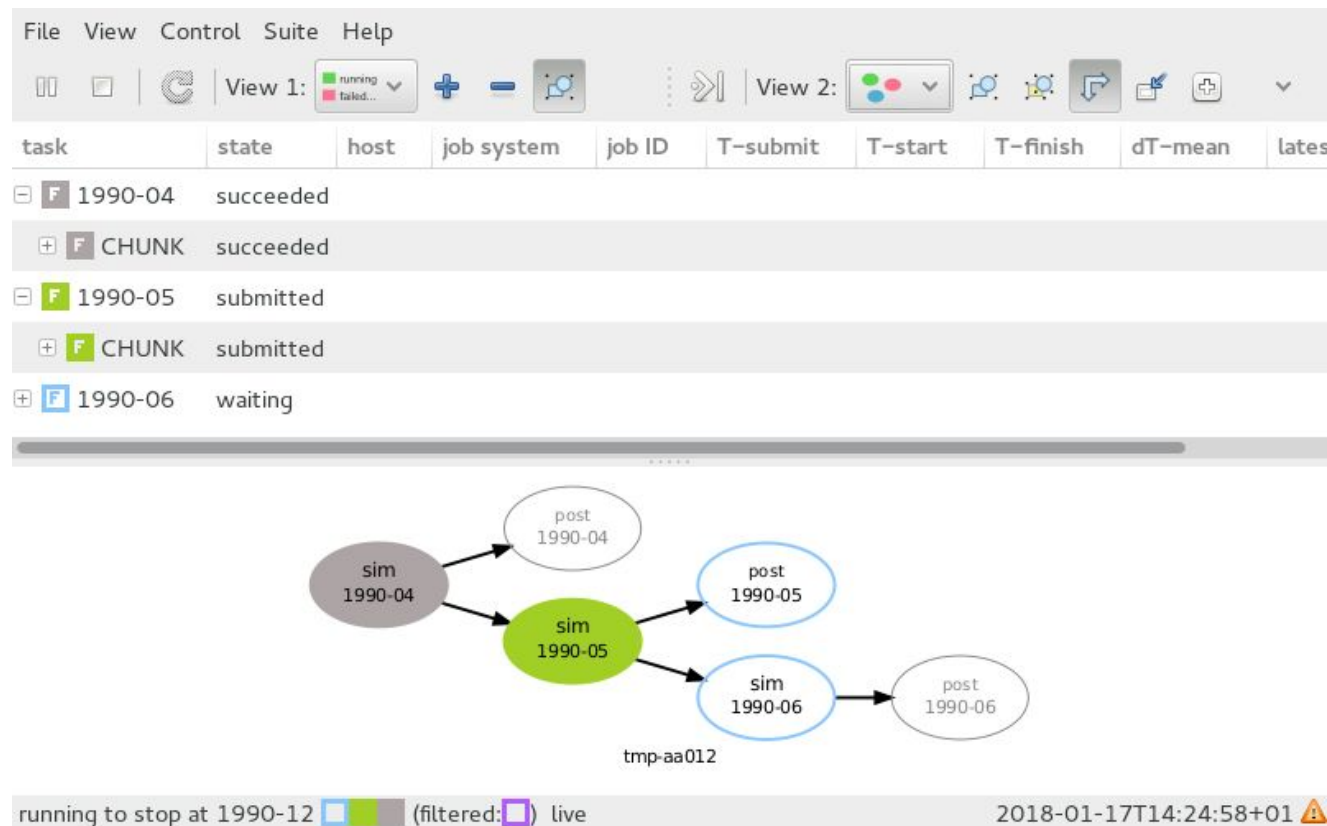
CMOR
standardized
(reduced)
outputs

Local storage



PA. Bretonnière (BSC)

- First **proof of concept** → NEMO standalone integrations
- “Recycled” Autosubmit EC-Earth **runtime**
- Testing different **interfaces** and **configurations**



File Edit View History Help

Search all revisions

roses:/

local	idx	owner	revision ^	title
🏠	tmp-aa012	user1	19	Test EC-Earth nemo only
🏠	tmp-aa011	user1	15	blah
🏠	tmp-aa010	dmanuben	12	spinup
📄	tmp-aa009	dmanuben	11	spinup
📄	tmp-aa008	dmanuben	10	spinup
🏠	tmp-aa007	dmanuben	9	spinup
🏠	tmp-aa006	dmanuben	8	spinup
🏠	tmp-aa004	dmanuben	6	spinup
🏠	tmp-aa003	dmanuben	5	tmp
🏠	tmp-aa002	dmanuben	4	tmp
🏠	tmp-aa001	dmanuben	3	tmp1
📄	tmp-aa000	dmanuben	2	test

tmp | 12 local suites found at 2018-01-17T14:02:28Z

Suites overview with Rosie Go

Suite configuration with Rose config edit

File Edit View Metadata Tools Page Help

Index

- suite info
- suite conf
- copy-runtime
 - command
- env
- post
 - command
 - env
- sim
 - command
 - env
- transfer
 - command
 - env

env X

Environment variable configuration

HPCARCH marenostum4

HPCHOST mn1.bsc.es

HPCROOTDIR /gpfs/scratch/bsc32704/tmp-aa012

HPCUSER bsc32704

Added env=SDATE_test



New developments

CMIP6 forcing



LCMIP6_STRAT_SIMP

“Simple scheme”
AOD at 550 nm,
vertically integrated

Default configuration
with LCMIP6=TRUE

LBCKGD_STRATAER
(without eruptions)

LCMIP6_STRAT_FULL

“Full scheme”
AOD vertically and
spectrally distributed

Configuration
available in the trunk,
Under validation

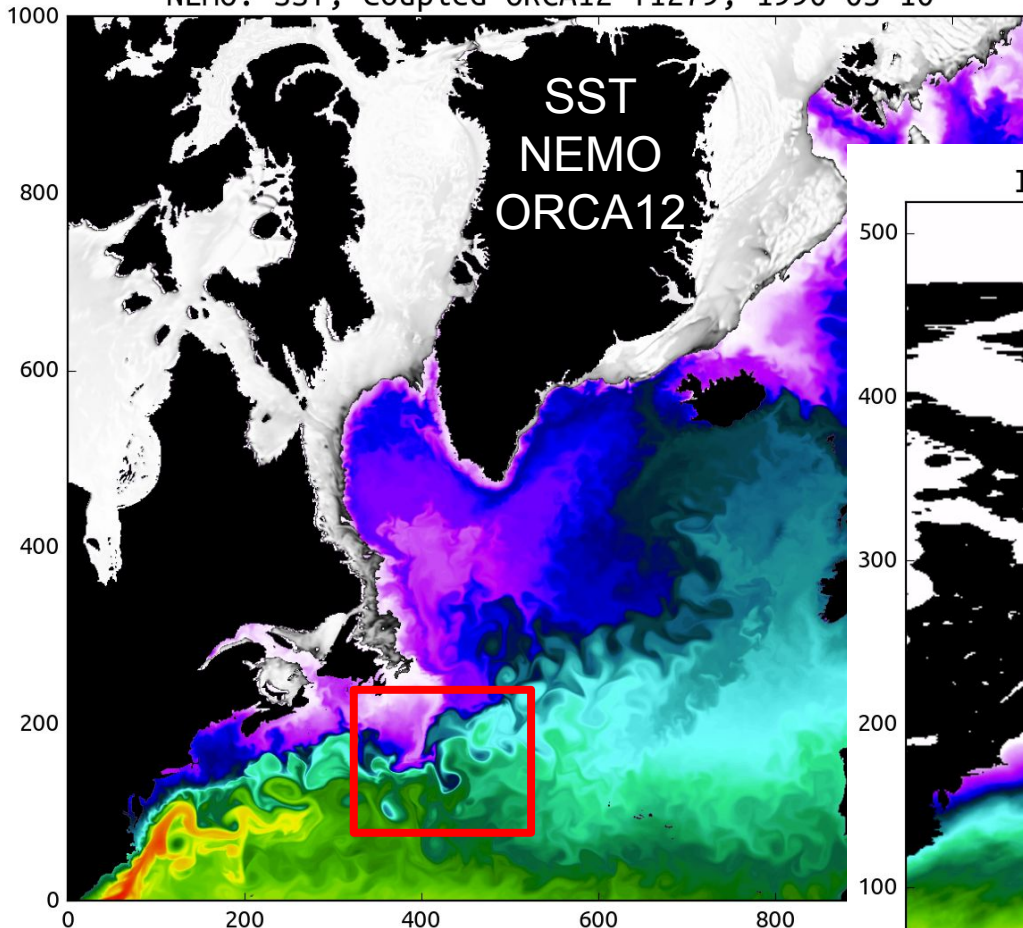
T1279-ORCA12



Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación




NEMO: SST, coupled ORCA12-T1279, 1990-03-10



EC-Earth GLOBAL ORCA12-T1279 (ocean and atmosphere at ~15 km!)

IFS: T2M, coupled ORCA12-T1279, 1990-03-10



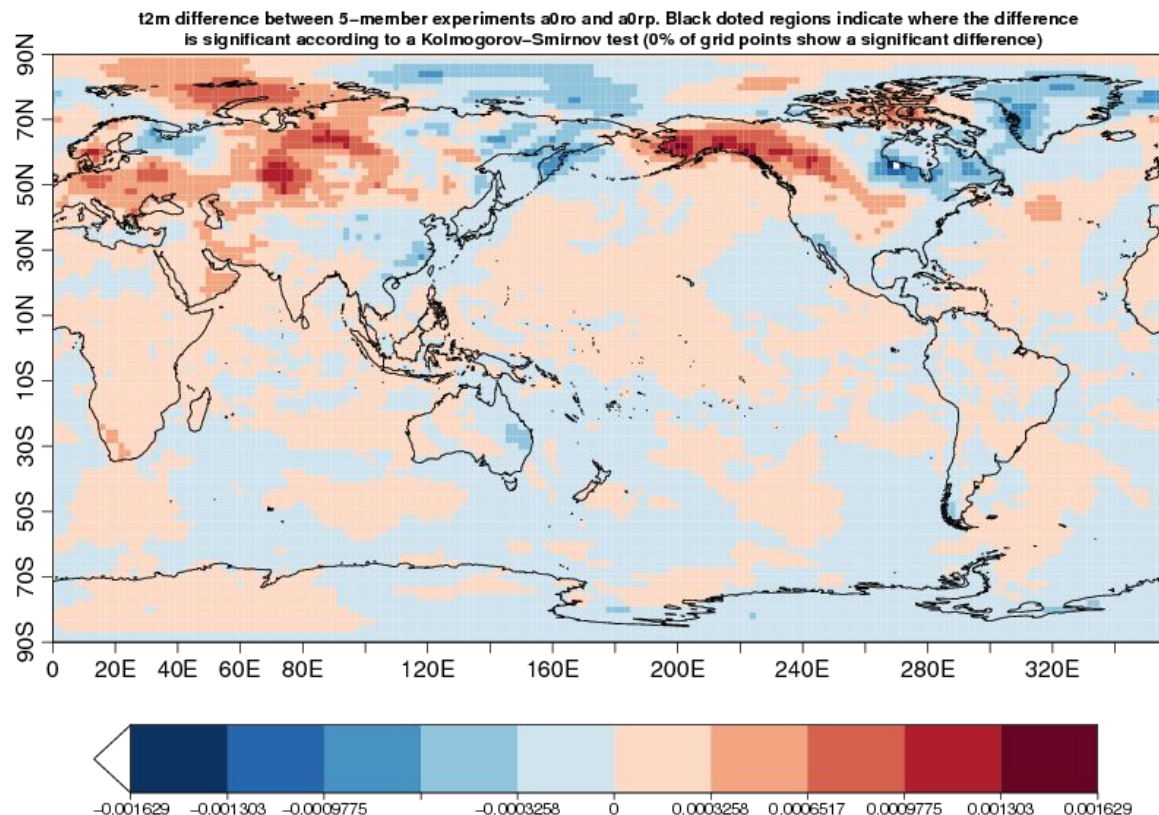
<https://www.youtube.com/watch?v=AU2zp7a8G0U>

laurent.brodeau@bsc.es

- Ocean **nudging** (surface and 3D) for **nemo-only** and **coupled**
 - V. Sicardi (BSC)
- Interannual variable **atmospheric forcings** management
 - V. Sicardi (BSC)
- **Pisces** fully **coupled** (IFS+NEMO+PISCES)
 - V. Sicardi (BSC)
- High resolution **sea-ice reanalysis** and **sea-ice starting conditions** for the period 1979-2016.
 - Sea ice concentration assimilation (Ensemble Kalman Filter method) - NEMO3.6-LIM3 (ORCA025)
 - J. Acosta (BSC), F. Massonet (UCL)

- **Initial conditions creation** during simulation
 - IFS & Nemo any time, Oasis only at start/end of leg run
 - E. Tourigny (BSC)
- **Land surface offline** script in ESM branch (ece-lsm.sh)
 - LPJ-Guess forced by daily IFS output
 - to be used for fire predictions using the output of DCPD experiments
 - OSM (offline surface model or H-TESSEL) has been updated by E. Dutra, integration to offline script to be done
 - E. Tourigny (BSC), P. Miller & L. Nieradzik (Lund), P. Anthoni (KIT), E. Dutra (U. Lisboa)

- The **methodology** is **ready** as starting point to evaluate reproducibility across platforms, compilation options, versions...
- Compare to CMIP5 results to evaluate the accuracy and between two experiments to evaluate the **reproducibility**
- Scripts are ready to use in a GitLab repository



M. Acosta (BSC),
F. Massonet (UCL),
M. Menegoz,
P. Le Sager (KNMI)

Thank you!



**Barcelona
Supercomputing
Center**

Centro Nacional de Supercomputación



EXCELENCIA
SEVERO
OCHOA

For further information please **contact**

Miguel Castrillo: miguel.castrillo@bsc.es

Pablo Echevarría: pablo.echevarria@bsc.es

Mario Acosta: mario.acosta@bsc.es

Domingo Manubens: domingo.manubens@bsc.es

Pierre-Antoine Bretonnière: pierre-antoine.bretonniere@bsc.es

Etienne Tourigny: etienne.tourigny@bsc.es

Oriol Tintó: oriol.tinto@bsc.es

Martin Menegoz: martin.menegoz@bsc.es

Valentina Sicardi: valentina.sicardi@bsc.es

Xavier Yepes: xavier.yepes@bsc.es

Juan Acosta: juan.acosta@bsc.es

Thomas Arsouze: thomas.arsouze@bsc.es